

ARM® Cortex™-A9 processors

r2 releases

Software Developers Errata Notice



ARM Cortex-A9 processors

Software Developers Errata Notice

Copyright © 2012 ARM Limited. All rights reserved.

Release Information

The following changes have been made to this book.

Change History

Date	Issue	Confidentiality	Change
18 June 2012	A	Non-confidential	First release for r2 releases, limited distribution
20 September 2012	B	Non-confidential	Second release for r2 releases

Proprietary Notice

This document is protected by copyright and the practice or implementation of the information herein may be protected by one or more patents or pending applications. No part of this document may be reproduced in any form by any means without the express prior written permission of ARM. No license, express or implied, by estoppel or otherwise to any intellectual property rights is granted by this document.

This document is Non-Confidential but any disclosure by you is subject to you providing the recipient the conditions set out in this notice and procuring the acceptance by the recipient of the conditions set out in this notice.

Your access to the information in this document is conditional upon your acceptance that you will not use, permit or procure others to use the information for the purposes of determining whether implementations infringe your rights or the rights of any third parties.

Unless otherwise stated in the terms of the Agreement, this document is provided "as is". ARM makes no representations or warranties, either express or implied, included but not limited to, warranties of merchantability, fitness for a particular purpose, or non-infringement, that the content of this document is suitable for any particular purpose or that any practice or implementation of the contents of the document will not infringe any third party patents, copyrights, trade secrets, or other rights. Further, ARM makes no representation with respect to, and has undertaken no analysis to identify or understand the scope and content of such third party patents, copyrights, trade secrets, or other rights.

This document may include technical inaccuracies or typographical errors.

TO THE EXTENT NOT PROHIBITED BY LAW, IN NO EVENT WILL ARM BE LIABLE FOR ANY DAMAGES, INCLUDING WITHOUT LIMITATION ANY DIRECT LOSS, LOST REVENUE, LOST PROFITS OR DATA, SPECIAL, INDIRECT, CONSEQUENTIAL, INCIDENTAL OR PUNITIVE DAMAGES, HOWEVER CAUSED AND REGARDLESS OF THE THEORY OF LIABILITY, ARISING OUT OF OR RELATED TO ANY FURNISHING, PRACTICING, MODIFYING OR ANY USE OF THIS DOCUMENT, EVEN IF ARM HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

Words and logos marked with ® or TM are registered trademarks or trademarks, respectively, of ARM Limited. Other brands and names mentioned herein may be the trademarks of their respective owners. Unless otherwise stated in the terms of the Agreement, you will not use or permit others to use any trademark of ARM Limited.

This document consists solely of commercial items. You shall be responsible for ensuring that any use, duplication or disclosure of this document complies fully with any relevant export laws and regulations to assure that this document or any portion thereof is not exported, directly or indirectly, in violation of such export laws.

In this document, where the term ARM is used to refer to the company it means "ARM or any of its subsidiaries as appropriate".

Copyright © 2012 ARM Limited

110 Fulbourn Road, Cambridge, England CB1 9NJ. All rights reserved.

Confidentiality Status

This document is Non-Confidential. The right to use, copy and disclose this document may be subject to license restrictions in accordance with the terms of the agreement entered into by ARM and the party that ARM delivered this document to.

Product Status

The information in this document is final, that is for a developed product.

Web Address

<http://www.arm.com>

Contents

ARM Cortex-A9 processors Software Developers

Errata Notice

Chapter 01	Introduction	
	1.1 Scope of this document	1-8
	1.2 Categorization of errata	1-9
	1.3 Errata summary	1-10
Chapter 02	Errata Descriptions	
	2.1 Category A	2-14
	2.2 Category A (Rare)	2-19
	2.3 Category B	2-26
	2.4 Category B (Rare)	2-47
	2.5 Category C	2-52

Chapter 1

Introduction

This chapter introduces the errata notices for ARM Cortex™-A9 processors, for r2 release.

1.1 Scope of this document

This document describes errata categorized by level of severity. Each description includes:

- The current status of the defect.
- Where the implementation deviates from the specification and the conditions under which erroneous behavior occurs.
- The implications of the erratum with respect to typical applications.
- The application and limitations of a work-around, where possible.

This document describes errata that may impact anyone who is developing software that will run on implementations of this ARM product.

1.2 Categorization of errata

Errata recorded in this document are split into the following levels of severity:

Table 1-1 Categorization of errata

Errata type	Definition
Category A	A critical error. No workaround is available or workarounds are impactful. The error is likely to be common for many systems and applications.
Category A (rare)	A critical error. No workaround is available or workarounds are impactful. The error is likely to be rare for most systems and applications. Rare is determined by analysis, verification and usage.
Category B	A significant error, or a critical error with an acceptable workaround. The error is likely to be common for many systems and applications.
Category B (rare)	A significant error, or a critical error with an acceptable workaround. The error is likely to be rare for most systems and applications. Rare is determined by analysis, verification and usage.
Category C	A minor error.

1.3 Errata summary

Table 1-2 lists all the errata described in this document. The Status column shows any errata that are new or updated in the current issue of the document. An erratum is shown as updated if there has been any change to the text of the erratum Description, Conditions, Implications or Workaround. Fixed errata are not shown as updated, unless the erratum text has changed.

Table 1-2 List of errata

Status	ID	Area	Cat	Rare	Summary of erratum
-	754319	Prog	A	-	A sequence of cancelled Advanced-SIMD or VFP stores might deadlock
-	754320	Prog	A	-	A cancelled Advanced-SIMD or VFP load multiple of more than 8 beats might deadlock
-	726781	Prog	A	Rare	Under very rare circumstances, an LDR/LDREX/STREX sequence might reach an unstable state, possibly leading to data corruption
-	742231	Prog	A	Rare	Incorrect hazard handling in the SCU might cause data corruption
-	761319	Prog	A	Rare	Possible inconsistent sequencing of read accesses to the same memory location
Updated	764269	Prog	A	Rare	Under very rare circumstances, a sequence of at least three writes merging in the same 64-bit address range might cause data corruption
-	745320	Prog	A	Rare	A Floating Point write following a failed conditional read might write corrupted data
-	729905	Prog	B	-	Using the PLE might cause a processor deadlock
-	740657	Prog	B	-	Global Timer can send two interrupts for the same event
-	742230	Prog	B	-	DMB operation might be faulty
-	751469	Prog	B	-	Overflow in PMU Counters might not be detected
-	751472	Prog	B	-	An interrupted ICIALUIS operation might prevent the completion of a following broadcast operation
-	751476	Prog	B	-	Missed watchpoint on the second part of an unaligned access crossing a page boundary
-	754322	Prog	B	-	Faulty MMU translations following ASID switch
Updated	764369	Prog	B	-	Data or unified cache line maintenance by MVA fails on Inner Shareable memory
-	775420	Prog	B	-	A data cache maintenance operation which aborts, followed by an ISB, without any DSB in-between may lead to deadlock
Updated	782772	Prog	B	-	Speculative execution of a Load-Exclusive or Store-Exclusive instruction after a write to Strongly Ordered memory might deadlock the processor
-	782773	Prog	B	-	Updating a translation entry to move a page mapping might erroneously cause an unexpected translation fault
New	794072	Prog	B	-	A short loop including a DMB instruction might cause a denial of service on another processor which executes a CP15 broadcast operation
New	794073	Prog	B	-	Speculative instruction fetches with MMU disabled might not comply with architectural requirements
New	794074	Prog	B	-	A write request to Uncacheable, Shareable normal memory region might be executed twice, possibly causing a software synchronisation issue

Table 1-2 List of errata (continued)

Status	ID	Area	Cat	Rare	Summary of erratum
-	743622	Prog	B	Rare	Faulty logic in the Store Buffer might cause data corruption
-	751473	Prog	B	Rare	Under very rare circumstances, automatic data prefetcher might cause deadlock or data corruption
Updated	761320	Prog	B	Rare	Full cache line writes to the same memory region from at least two processors might deadlock the processor
-	725631	Prog	C	-	ISB is counted in Performance Monitor events 0x0C and 0x0D
-	729808	Prog	C	-	PLE Stride ranges from 0 to 255, instead of from 1 to 256
-	729814	Prog	C	-	CFGSDISABLE pin value modification might not be detected
-	729816	Prog	C	-	TLB Lockdown write entry operation might not execute
-	729817	Prog	C	-	Main ID register alias addresses are not mapped on Debug APB interface
-	729818	Prog	C	-	In debug state, next instruction is stalled when sdabort flag is set, instead of being discarded
-	730345	Prog	C	-	PLE Wait States resolution is incorrect
-	740661	Prog	C	-	Event 0x74 / PMUEVENT[38:37] might be inaccurate
-	740663	Prog	C	-	Event 0x68 / PMUEVENT[9:8] might be inaccurate
-	743623	Prog	C	-	Bad interaction between a minimum of seven PLDs and one Non-Cacheable LDM, can lead to deadlock
-	743625	Prog	C	-	A coherent ACP request might interfere with a non-cacheable SWP/SWPB from the processor, potentially causing deadlock
-	743626	Prog	C	-	An imprecise external abort received while the processor enters WFI may cause a processor deadlock
-	751471	Prog	C	-	DBGPCSR format is incorrect
-	751480	Prog	C	-	Conditional failed LDREXcc can set the exclusive monitor
-	752519	Prog	C	-	An imprecise abort might be reported twice on non-cacheable reads
-	754323	Prog	C	-	Repeated Store in the same cache line might delay the visibility of the Store
-	756421	Prog	C	-	Sticky Pipeline Advance bit cannot be cleared from debug APB accesses
-	757119	Prog	C	-	Some <i>Unallocated memory hint</i> instructions generate an Undefined Instruction exception instead of being treated as NOP
-	761321	Prog	C	-	MRC and MCR are not counted
-	764319	Prog	C	-	Read accesses to DBGPRSR and DBGOSLSR may generate an unexpected Undefined Instruction exception
-	771221	Prog	C	-	PLD instructions might allocate data in the Data Cache regardless of the Cache Enable bit value
-	771224	Prog	C	-	Visibility of Debug Enable access rights to enable/disable tracing is not ensured by an ISB

Table 1-2 List of errata (continued)

Status	ID	Area	Cat	Rare	Summary of erratum
-	771225	Prog	C	-	Speculative cacheable reads to aborting memory regions clear the internal exclusive monitor, may lead to livelock
-	775419	Prog	C	-	PMU event 0x0A (exception return) might count twice the LDM PC ^ instructions with base address register write-back
-	782774	Prog	C	-	A spurious event 0x63, STREX passed, can be reported on an LDREX that is preceded by a write to Strongly Ordered memory region.

Chapter 2

Errata Descriptions

This chapter includes the errata descriptions for ARM Cortex™-A9 processors, for r2 releases.

2.1 Category A

This section describes Category A errata.

2.1.1 (754319) A sequence of cancelled Advanced-SIMD or VFP stores might deadlock

Status

Affects: Product Cortex-A9, Cortex-A9 MPCore

Fault Type: Programmer Category A

Fault Status: Present in: All r0, r1, r2 and r3 revisions. Fixed in r4p0, and FPU/MPE revision 0x4.

For all revisions, an ECO fix is available in the MPE logic. All designs having implemented the ECO are identifiable by reading FPSID.Revision=0x4.

Description

A store is *cancelled* in the following circumstances:

- if it fails its condition-code
- if it generates a precise abort
- if it forms part of speculative execution which is subsequently discarded as a result of the processor identifying a branch miss-predict scenario.

A store is *queued* after the integer core has determined the address of the transaction. A store can optionally not be queued if the integer core can determine that the store will be cancelled while the address computation is being performed, for example if the store is conditional and the CPSR flags are already ready.

A store is considered *ready* when the values of the operands to be stored have been computed. The order in which a sequence of stores become ready can differ from program order if the operation required to compute a result for an earlier store takes longer than that of the result required for a store that occurs later in the program order.

A store is considered *committed* after it is ready and can no longer be cancelled; in other words, after the value to be stored is known, the instruction passes any condition-code check, and any branch-prediction is known to be correct.

An Advanced-SIMD (Neon) store, or a VFP store executed as part of a Neon code sequence, can cause deadlock in the following circumstances:

1. The sequence results in the first store in the instruction order becoming ready and committed later than a set of following stores.
2. The subsequent stores become queued so that the second store is cancelled, a third store is committed, and a final store is cancelled.

Conditions

Group A is the set of Advanced-SIMD (Neon) or VFP stores whose data fits within a single 64-bit aligned 64-bit memory location, and which are issued by the Neon unit in a single cycle. This group contains all addressing mode forms of:

1. Single-precision 32-bit VFP stores:
VSTR.32 St, [Rn, ...]
2. Neon single D-register structure stores with alignment specifiers equal to 64-bits:
VST1.8 {Dd}, [Rn@64]
VST1.16 {Dd}, [Rn@64]
VST1.32 {Dd}, [Rn@64]

VST1.64 {Dd}, [Rn@64]

3. Neon indexed single 8-bit element store from one lane:
VST1.8 {Dd[<index>]}, [Rn]
4. Neon indexed single larger than 8-bit element store from one lane with an alignment specifier:
VST1.16 {Dd[<index>]}, [Rn@16]
VST1.32 {Dd[<index>]}, [Rn@32]
5. Neon indexed multi-element store of no more than 8-bytes with an alignment specifier:
VST2.8 {Dd[<index>], Dd2[<index>]}, [Rn@16]
VST2.16 {Dd[<index>], Dd2[<index>]}, [Rn@32]
VST2.32 {Dd[<index>], Dd2[<index>]}, [Rn@64]
VST4.8 {Dd[<index>], Dd2[<index>], Dd3[<index>], Dd4[<index>]}, [Rn@32]
VST4.16 {Dd[<index>], Dd2[<index>], Dd3[<index>], Dd4[<index>]}, [Rn@64]

The conditions for the erratum arise only when the following sequence occurs:

1. The Advanced-SIMD (Neon) extension is present, enabled and active (see below).
2. A group-A Neon or VFP store is queued, but remains not ready.
3. A Neon or VFP store is queued and cancelled, occupying only one 64-bit slot.
4. A group-A Neon or VFP store is queued and committed.
5. A Neon or VFP store is queued and cancelled.
6. The Neon or VFP store in (2) becomes ready and is committed.

Neon is active if an Advanced SIMD (Neon) instruction has been decoded more recently (including speculatively) than an instruction from any of the following categories:

1. A VFP data-processing instruction
2. A VFP-to-VFP register move (including moving a register to itself as a NOP)
3. A VFP condition code transfer (VMRS APSR_nzcv, FPSCR, that is FMSTAT)
4. A transfer (VMSR) from an integer register to a VFP system register
5. FLDMX/FSTMX (deprecated)

A VFP load/store instruction, or move between VFP and integer registers, has no effect on whether Neon is active. In particular, VFP/Neon register save/restore sequences typically occurring in context switches have no effect on whether Neon is active, unless they use the deprecated FSTMX/FLDMX forms.

Neon may be known to be inactive if there have been no Neon instructions decoded (including on any possible speculative paths) since the most recent instruction from one of the above VFP categories.

As an example:

```
LDR      r5, [r4]
CMP      r5, r5          ; delayed update of CPSR flags
VMUL.U32 d0, d0, d0      ; slow operation updating d0 (s0 and s1)
VSTR     s0, [r0]
...
VSTRNE   s2, [r1]        ; cancelled
...
```

VSTR s3,[r2] ; committed

...

VSTRNE s4,[r3] ; cancelled

Note that it is not a requirement for the Neon or VFP store instructions to be consecutive in the program. It is possible to separate these using a number of ARM register based instructions, branches or any other sequence not provided in the workarounds below. In addition the cancelled stores can reside in speculated branch shadows, and might not form part of a simple sequential execution of the same code.

This erratum is timing sensitive and is influenced by the relative cycle timings of the CPSR flags, branch prediction results, loads, stores, and both Neon and non-Neon related instructions becoming ready, all of which might not be predictable.

Implications

Only implementations of Cortex-A9, including the Multimedia-Processing Engine (MPE) / Neon unit, are affected by this erratum. Execution of a code sequence that stimulates this erratum ultimately results in deadlock. Neon stores that are committed after the erratum is triggered, but before deadlock occurs, might use the correct address but store out-of-sequence Neon data.

Workaround

This erratum does not affect implementations that do not have Neon, or implementations that operate with Neon disabled. You can use the ASEDIS bit to disable the Neon functionality and avoid this erratum, while still retaining VFP floating-point capabilities.

You can enable software that uses Neon to work around this erratum by inserting any of the following:

- One or more Neon or floating-point instructions between conditions 2 and 3.
- Two or more Neon or floating-point instructions between conditions 3 and 4.
- One or more Neon or floating-point load instructions between conditions 3 and 4.
- One or more Neon or floating-point instructions between conditions 4 and 5.

When performing code insertion, take care not to re-create an alternative sequence susceptible to this erratum.

You can enable software using Neon to work around this erratum by substituting the instructions in Group A with those not from Group A as follows:

1. You can correct instructions that contain alignment specifiers by substitution with the identical instruction minus the alignment specifier.
2. VSTR instructions with zero immediate offset may be code substituted a single word VSTM, for example, VSTR s3,[r0] can be replaced with VSTM r0,{s3}.

————— Note —————

Note that this code substitution does not provide equivalents for single-lane variant of VST1.8 or for VSTR instructions with non-zero offsets. For these instructions, code insertion remains the preferred corrective action.

Suitable instructions for code insertion include VORR d0,d0,d0 for Neon intensive code, and VMOV.F32 s0,s0 for VFP intensive code. Alternatively, you can insert an additional VSTR.F64 to a scratch location.

Alternatively, software that does not use Neon can avoid the need to apply the workaround by ensuring that it does not enter the code (including by way of an exception return) with Neon active. For exception returns to software that is not using Neon, you can ensure this by having a VMOV.F32 s0,s0, or other VFP instruction, that prevents Neon being active (as described above) in the final basic block before the exception return, and by ensuring that the bit pattern immediately following the exception return instruction does not correspond to an instruction that makes Neon active.

2.1.2 (754320) A cancelled Advanced-SIMD or VFP load multiple of more than 8 beats might deadlock

Status

Affects: Product Cortex-A9, Cortex-A9 MPCore

Fault Type: Programmer Category A

Fault Status: Present in: All r0, r1, r2 and r3 revisions. Fixed in r4p0, and FPU/MPE revision 0x4.

For r0, r1, r2 and r3 revisions, an ECO fix is available in the MPE logic. All designs having implemented the ECO are identifiable by reading FPSID.Revision=0x4.

Description

A *beat* is required for each 64-bit aligned chunk of 64-bits required to service a load. This erratum affects VLDM and VPOP instructions that require more than 8 beats only. This includes all Advanced-SIMD or VFP load multiple instructions which load:

- greater-than 8 double-word D-register or 16 single-word S-registers from a 64-bit aligned address
- greater-than 7 double-word D-register or 15 single-word S-registers from a less than 64-bit aligned address.

A load is *cancelled* if it fails its condition code, or if it forms part of speculative execution which is subsequently discarded because the processor identifies a branch miss-predict scenario.

A speculative Advanced-SIMD or VFP load multiple of more than eight beats speculatively executed by the Neon unit executed as part of a Neon code sequence, but subsequently cancelled because of miss-speculation, can result in deadlock if followed by a committed Advanced-SIMD or VFP load.

Conditions

1. The Advanced-SIMD (Neon) extension is present and enabled.
2. A Neon or VFP load multiple of more than 8-beats is speculatively issued
3. A Neon or VFP load is issued and not cancelled.
4. The Neon or VFP load in (2) is cancelled.

As an example:

```
CMP    r0,r0; ensure EQ condition passes
BEQ    1      ; conditional branch to VLDR
VLDM    r0,{d0-d9} ; cancelled due to branch miss-predict
...
1:      VLDR    s0,[r1] ; executed and not cancelled
```

This erratum is timing sensitive and is influenced by the relative cycle timings of the CPSR flags, branch prediction results, loads, stores, and the Neon, VFP and integer instructions becoming ready, all of which might not be predictable.

Implications

This erratum affects only implementations of Cortex-A9 that include the Multimedia-Processing-Engine (MPE) or Neon unit. Execution of a code sequence that stimulates this erratum might result in deadlock.

Workaround

This erratum does not affect implementations without Neon, or those that operate with Neon disabled. You can use the ASEDIS bit to disable the Neon functionality and avoid this erratum, while still retaining VFP floating-point capabilities.

Software that operates with the Advanced-SIMD (Neon) extension enabled and executing Neon software can work around this erratum by splitting load multiples that are capable of generating more than 8-beats into smaller loads.

Note that compliance to the AAPCS ensures that, on exit from a public interface, this erratum will not affect a callee restore of the form `VP0P {d8-d15}` because of the AAPCS requirement for 64-bit stack-alignment.

2.2 Category A (Rare)

This section describes Category A rare errata.

2.2.1 (726781) Under very rare circumstances, an LDR/LDREX/STREX sequence might reach an unstable state, possibly leading to data corruption

Status

Affects: Product Cortex-A9 MPCore.

Fault Type: Programmer Category A (Rare)

Fault Status: Present in: r2p0. Fixed in r2p1.

Description

Under very rare circumstances, an LDR/LDREX/STREX sequence can reach an unstable state, possibly leading to data corruption.

Conditions

- Cortex-A9 MPCore configuration with coherency logic present, that is:
 - one CPU and ACP present, and coherent requests occur on ACP
 - two CPUs or more
- Data Cache on
- CPU is part of the coherency domain (ACTLR.SMP=1).

The CPU then needs to execute the following code sequence, possibly with other instructions interleaved:

- LDR in a cacheable memory region, that misses in the L1 Data Cache
- LDREX in the same cache line
- STREX in the same cache line.

The failure also requires the failing CPU to receive a coherent request.

With the code sequence above, under certain timing conditions specific to the Cortex-A9 micro-architecture, the Cortex-A9 might enter an unexpected state, which might cause any future STREX to cause a data corruption.

Implications

The erratum possibly leads to data corruption.

Workaround

The only reliable software workaround for this erratum is to insert a DSB between the LDR and the LDREX.

Because the DSB is a strong memory barrier, which waits for all Loads and Stores in the Cortex-A9 to complete, this workaround can significantly degrade the performance of the LDR / LDREX / STREX sequence.

2.2.2 (742231) Incorrect hazard handling in the SCU might cause data corruption

Status

Affects: Product Cortex-A9 MPCore.

Fault Type: Programmer Category A (Rare)

Fault Status: Present in: r2p0, r2p1, r2p2 Fixed in r2p5, r2p6

Description

Under very rare circumstances, a cache line shared between two coherent agents (for example, two different processors in the Cortex-A9 MPCore cluster, or one processor and the ACP) might be corrupted.

Conditions

The erratum occurs only in Cortex-A9 MPCore configurations with coherency logic, that is, those containing one processor with the ACP present, or those containing two or more processors (with or without ACP).

The two coherent agents must work in SMP mode, on the same cache line. Under very rare timing conditions, specific to the Cortex-A9 micro-architecture, the cache line which is shared between the two agents might be corrupted because of a faulty hazard checking mechanism in the SCU.

In SMP mode, the SCU hazard logic might not trigger a hazard between a line currently in the eviction buffer of a first processor, and a coherency request from another coherent master.

The result is that the memory request is serviced from the next memory level, resulting in the loss of the data currently being evicted.

ARM has not observed this erratum in systems that use the workaround to erratum 742230. However, ARM cannot guarantee freedom from this erratum in some memory systems with some particular timing.

Implications

This erratum might cause data corruption.

Workaround

For Cortex-A9 MPCore configurations containing two or more processors, the software workaround is to set bit[12] in the undocumented Diagnostic Control register placed in CP15 c15 0 c0 1.

It is possible to write this bit in Secure state only, with the following Read/Modify/Write code sequence:

```
MRC      p15,0,rt,c15,c0,1
ORR      rt,rt,#0x1000
MCR      p15,0,rt,c15,c0,1
```

Bit[12] is the Disable Outstanding Accesses bit, which prevents the processor performing more than two data line-fill requests at a time.

When you use this workaround, there is likely to be a visible drop in performance on routines that perform intensive memory accesses, such as a `memcpy()`. However, the workaround is not likely to create any significant performance degradation in other standard applications.

For Cortex-A9 MPCore configuration with a single processor and the ACP present, the software workaround is slightly different, and involves settings bit[12] and bit[22] in the undocumented Diagnostic Control register.

It is possible to write these two bits in Secure state only, with the following Read/Modify/Write code sequence:

```
MRC      p15,0,rt,c15,c0,1
```

```
ORR      rt,rt,#0x00001000
ORR      rt,rt,#0x00400000
MCR      p15,0,rt,c15,c0,1
```

Bit[12] is the Disable Outstanding Accesses bit, which prevents the processor performing more than two data line-fill requests at a time.

Bit[22] is the Disable Wait Mode bit, which prevents automatic switching of the processor in Read-Allocate mode when identifying patterns such as a `memset()` or `memcpy()` routine.

ARM recommends using the hardware ECO to fix any configuration affected by this erratum. Note that the ECO also fixes erratum #742230. The hardware ECO is available as a patch release for r2p1 and r2p2 revisions of the product.

- r2p1-50rel0 contains the ECO patch for the r2p1 revision of the product, and upgrade it to r2p5.
- r2p2-50rel0 contains the ECO patch for the r2p2 revision of the product, and upgrade it to r2p6.

2.2.3 (745320) A Floating Point write following a failed conditional read might write corrupted data

Status

Affects: Product Cortex-A9, Cortex-A9 MPCore.

Fault Type: Programmer Category A (Rare)

Fault Status: Present in: All r0, r1 and r2 revisions* Fixed in r3p0, and FPU/MPE revision 0x3

* For r0, r1 and r2 revisions, an ECO fix is available in the FPU or MPE logic. You can read FPSID.Revision=0x3 to identify all designs that have implemented the ECO.

Description

Under certain conditions specific to the Cortex-A9 micro-architecture, a floating point write operation which executes a few cycles after a failed conditional read might write corrupted data to the destination address.

Conditions

The erratum is present only in processor configurations where a Data Engine is present (either the FPU, or the Neon SIMD engine which automatically includes the FPU).

The processor needs to execute a conditional load instruction which fails its condition check.

An FPU write operation occurs (VSTR/VSTM/VPUSH/FSTMX) within a few cycles following the conditional load instruction.

Under specific timing conditions, the failed conditional read might interfere with the floating point write, possibly causing a data corruption for the written data.

Implications

The erratum might cause a data corruption on the floating point data.

Workaround

For manual programming, or if the compiler offers the capability, a possible software workaround is to prevent the use of conditional loads in programs which make use of the floating point unit.

There is no practical software workaround if it is not possible to avoid the offending sequence.

2.2.4 (761319) Ordering of read accesses to the same memory location might be uncertain

Status

Affects: Product Cortex-A9 MPCore

Fault Type: Programmer Category A (Rare)

Fault Status: Present in: All revisions. Open

Description

The ARM architecture and the general rules of coherency require reads to the same memory location to be observed in sequential order.

Because of some internal replay path mechanisms, the Cortex-A9 can see one read access bypassed by a following read access to the same memory location, thus not observing the values in program order.

Conditions

The erratum requires a Cortex-A9 MPCore configuration with two or more processors or more.

The erratum can occur only on a processor working in SMP mode, on memory regions marked as Normal Memory Write-Back Shared.

Implications

The erratum causes data coherency failure.

Workaround

The majority of multi-processing code examples follow styles that do not expose the erratum. Therefore, this erratum occurs rarely and is likely to affect only very specific areas of code that rely on a read-ordering behavior.

There are two possible workarounds for this erratum:

- The first possible workaround is to use LDREX instead of standard LDR in volatile memory places that require a strict read ordering.
- The alternative possible workaround is the recommended workaround for tool chains integration. It requires insertion of a DMB between the affected LDR that requires this strict ordering rule.

For more information about integrating the workaround inside tool chains, see the Programmer Advice Notice related to this erratum, *ARM UAN 0004A*.

2.2.5 (764269) Under very rare circumstances, a sequence of at least three writes merging in the same 64-bit address range might cause data corruption

Status

Affects: Product Cortex-A9, Cortex-A9 MPCore.

Fault Type: Programmer Category A (Rare)

Fault Status: Present in: All r2 revisions. Fixed in r3p0

Description

Under very rare timing circumstances, specific to the Cortex-A9 micro-architecture, a sequence involving at least three writes merging into the same 64-bit cacheable memory region, might cause data corruption.

Conditions

The erratum requires the following:

- A cacheable write at address A, in a cache line which is currently being brought in the processor because of a previous request.
- A write at another address B, in a different cache line from A. This write must also miss its cache lookup inside the data cache.
- A second write in the same naturally-aligned 64-bit region as address A. This second write must merge in the same 64-bit slot as address A in the clock cycle when this slot is either merging, or checking its merging status, in the Bus Interface Unit.
- The write at address B which had missed should take priority over the second write at address A which is being replayed in the Bus Interface Unit.
- A cache line, different from the one containing address A, must be allocated in the following cycle.
- A third write in the same 64-bit region as address A. This third write must merge in the same 64-bit slot as the first write at address A. This merge must occur in exactly the same cycle as when cache line A is allocated in the data cache.
- When the write A is replayed on the data cache, it must miss.

The code sequence generating these accesses must be seen by the Cortex-A9 Store Buffer in a short timing window of approximately 11 cycles. Note that this code sequence can include additional instructions, for example another write in cache line B between the second and third writes to A.

Because the erratum requires three writes which merge in the same 64-bit region in a very short period, it is more likely to be triggered in a code sequence involving byte or half-word stores, because three word stores would expand beyond the mandatory 64-bit region.

As an example, the following loop is theoretically susceptible to the erratum:

```
loop    LDRB A
        LDRB B
        STRB A
        STRB B
        SUBS
        BNE    loop
```


As this description shows, the erratum requires very rare and precise alignment of different micro-architectural events. These events are not controllable by the software, but this requirement explains why this example code sequence executes correctly in most cases.

In practice, the erratum is likely to happen in exceptionally rare conditions. ARM has not managed to reproduce the failure in a top-level simulation because of the extreme difficulty in setting up the timing conditions of the sequence of micro-architectural events that are necessary for the failure to occur.

Implications

When the erratum occurs, the allocation of the written 64-bit data slot at address A is wrongly cancelled, causing data corruption.

Workaround

There is no practical workaround for this erratum.

2.3 Category B

This section describes Category B errata.

2.3.1 (729905) Using the PLE might cause a processor deadlock

Status

Affects: Product Cortex-A9, Cortex-A9 MPCore.

Fault Type: Programmer Category B

Fault Status: Present in: r2p0,r2p1,r2p5,r2p7,r2p9 Fixed in all others

Description

When multiple requests generated by the PLE are pending in the processor, the hazard checking mechanism between the different requests might become corrupted, possibly leading to a processor deadlock.

Conditions

Requires a hardware configuration which includes the Preload Engine (PLE).

Implications

Under some precise timing conditions specific to the Cortex-A9 micro-architecture, when the PLE generates multiple outstanding requests to the external memory system, a deadlock can occur caused by faulty hazard checking.

The erratum is more likely to happen in configurations with a slow memory system, or with a high PLE issuing rate.

Workaround

The only reliable workaround for this erratum is to not use the PLE feature in the revisions of the product which are affected by this erratum.

Secure code can prevent the Non-Secure world using the PLE by maintaining NSACR.PLE=0, which is the default value.

Privilege modes can also prevent User code programming any PLE by maintaining PLEUAR.U=0, which is the default value.

2.3.2 (740657) Global Timer can send two interrupts for the same event

Status

Affects: Product Cortex-A9 MPCore.

Fault Type: Programmer Category B

Fault Status: Present in: All r2, r3 and r4 revisions. Open

Description

The Global Timer can be programmed to generate an interrupt request to the processor when it reaches a given programmed value. Because of the erratum, when you program the Global Timer to not use the auto-increment feature, it might generate two interrupt requests instead of one.

Conditions

The Global Timer Control register is programmed with the following settings:

- Bit[3] = 1'b0 - Global Timer is programmed in single-shot mode
- Bit[2] = 1'b1 - Global Timer IRQ generation is enabled
- Bit[1] = 1'b1 - Global Timer value comparison with Comparator registers is enabled
- Bit[0] = 1'b1 - Global Timer count is enabled.

With these settings, an IRQ is generated to the processor when the Global Timer value reaches the value programmed in the Comparator registers. The Interrupt Handler then performs the following sequence:

- Read the ICCIAR (Interrupt Acknowledge) register
- Clear the Global Timer flag
- Modify the comparator value, to set it to a higher value
- Write the ICCEOIR (End of Interrupt) register.

Under these conditions, because of the erratum the Global Timer might generate a second (spurious) interrupt request to the processor at the end of this Interrupt Handler sequence.

Implications

The erratum creates spurious interrupt requests in the system.

Workaround

Because the erratum happens only when the Global Timer is programmed in single-shot mode, that is, when it does not use the auto-increment feature, a first possible workaround is to program the Global Timer to use the auto-increment feature.

If this first solution is not possible, a second workaround is to modify the Interrupt Handler to avoid the offending sequence. You can achieve this by clearing the Global Timer flag after incrementing the Comparator register value. The correct code sequence for the Interrupt Handler should then look like the following sequence:

- Read the ICCIAR (Interrupt Acknowledge) register
- Modify the comparator value, to set it to a higher value
- Clear the Global Timer flag

- Clear the Pending Status information for Interrupt 27 (Global Timer interrupt) in the Distributor of the Interrupt Controller.
- Write the ICCEOIR (End of Interrupt) register.

2.3.3 (742230) DMB operation might be faulty

Status

Affects: Product Cortex-A9 MPCore.

Fault Type: Programmer Category B

Fault Status: Present in: r1p0, r1p1, r1p2, r1p3, r2p0, r2p1, r2p2 Fixed in all other revisions.

Description

On Cortex-A9, the Data Memory Barrier operation might exhibit a faulty behavior by failing to ensure that no reordering occurs between two write operations separated by a DMB.

Conditions

The erratum occurs only in Cortex-A9 MPCore configurations with coherency logic, that is, those containing one processor with the ACP present, or those containing two or more processors with or without ACP.

Under rare circumstances, specific to the Cortex-A9 MPCore micro-architecture, the DMB operation might complete too early, allowing a subsequent write instruction to proceed and to be made visible to external agents while a previous write instruction that occurred before the DMB in program order might not have completed. The result of this write is not yet visible to the external agent.

The processor must perform:

- A write to a first address, which misses in the L1 Data Cache, and causes a LFI (Line-fill and Invalidate) request to the SCU
- A DMB
- A write to a second address, located on a different cache line, which hits in the L1 Data Cache.

To observe the memory ordering problem and create the errata the external agent - either another processor in the MPCore cluster or an agent connected to the ACP, must:

- Issue a memory request to the second address, which should cache miss and cause a coherent line fill request to the SCU. The SCU must arbitrate and service this coherent line-fill request before the LFI request from the first processor. The external agent then receives the data written by the second write of the first processor.
- Perform a read or a write memory request to the address targeted by the first write of the first processor, and this read must cache hit so that it gets the old data instead of the data written by the first processor.

Note that the erratum occurs independently of the Shareability domain of the DMB operation.

Implications

Because of the erratum, an external agent might observe the result of an explicit write instruction that occurred later than the DMB instruction in program order before it observes the results of an explicit write instruction that occurred earlier than the DMB instruction in program order.

Workaround

The software workaround for this erratum is to set bit[4] in the undocumented Diagnostic Control register placed in CP15 c15 0 c0 1.

You can write this bit in Secure state only, using the following Read/Modify/Write code sequence:

```
MRC p15,0,rt,c15,c0,1
ORR rt,rt,#0x10
MCR p15,0,rt,c15,c0,1
```

When this bit is set, it causes the DMB instruction to be decoded and executed like a DSB in the processor.

This software workaround has no expected visible impact on the overall performance of the processor on a typical code base.

The ECO applied to fix erratum 742231 (see [page 2-20](#)) also addresses this erratum.

2.3.4 (751469) Overflow in PMU Counters might not be detected

Status

Affects: Product Cortex-A9, Cortex-A9 MPCore

Fault Type: Programmer Category B

Fault Status: Present in: All r0, r1 and r2 revisions, Fixed in r3p0

Description

Overflow detection logic in the Performance Monitor Counters is faulty, and under certain timing conditions the overflow might remain undetected. In this case, the Overflow Flag Status register (PMOVSr) does not update as it should, and there is no interrupt reported on the corresponding PMUIRQ line.

Implications

PMU overflow detection is not reliable.

Workaround

The workaround for this erratum is to set two PMU counters to count the same event, and to explicitly offset them by 1 at the start of the count. The following sequence achieves this:

1. Disable PMU count.
2. Set up Counter0 to value N.
3. Set up Counter1 to value (N+1).
4. Enable PMU count.

This ensures that at least one of the two counters detects the overflow.

Usually, both counters trigger. Therefore, if using this workaround it is necessary for the software to reset both counters when the first one triggers.

2.3.5 (751472) An interrupted ICIALLUIS operation might prevent the completion of a following broadcast operation

Status

Affects: Product Cortex-A9 MPCore

Fault Type: Programmer Category B

Fault Status: Present in: All r0, r1 and r2 revisions. Fixed in r3p0

Description

In an MPCore configuration with two or more processors working in SMP mode with maintenance operation broadcast enabled, if a processor is interrupted while executing an ICIALLUIS operation, and performs another broadcast maintenance operation during its Interrupt Service Routine, then this second operation might not execute on the other processors in the cluster.

Conditions

The erratum requires an MPCore configuration with two or more CPUs working in SMP mode.

One processor has interrupts, Cache, and TLB maintenance broadcast enabled (ACTLR.FW=1'b1). This processor executes an ICIALLUIS (Invalidate All Instruction Caches Inner Shareable to Point of Unification). This instruction executes on the processor, and is broadcast to other processors in the MPCore cluster.

The processor then receives an interrupt (IRQ or FIQ), which interrupts the ICIALLUIS operation.

During the Interrupt Service Routine, the processor executes any other Cache or TLB maintenance operation which is also broadcast to other processors in the MPCore cluster.

The erratum occurs if the other processors in the cluster receive this second maintenance operation before they complete the first ICIALLUIS operation, because the other processors do not execute the second maintenance operation. This is because there is no stacking mechanism for acknowledge answers between the processors, so that the originating processor interprets the acknowledge request sent to signify the completion of the ICIALLUIS as an acknowledgement for the second maintenance operation.

Implications

Because of the erratum, the processor might hold corrupted entries in the Cache or in the TLB, causing indeterminate failures in the system.

Workaround

A software workaround for this erratum is to set bit[11] in the undocumented Diagnostic Control register placed in CP15 c15 0 c0 1.

You can write this bit in Secure state only, by using the following Read/Modify/Write code sequence:

```
MRC p15,0,rt,c15,c0,1
```

```
ORR rt,rt,#0x800
```

```
MCR p15,0,rt,c15,c0,1
```

When it is set, this bit prevents the interruption of CP15 maintenance operations.

There is unlikely to be any visible impact on system performance when using this software workaround.

2.3.6 (751476) Missed watchpoint on the second part of an unaligned access crossing a page boundary

Status

Affects: Product Cortex-A9, Cortex-A9 MPCore

Fault Type: Programmer Category B

Fault Status: Present in: All revisions. Open

Description

Under rare conditions, a watchpoint might be undetected if it occurs on the second part of an unaligned access that crosses a 4K page boundary and misses in the μ TLB for the second part of its request.

The erratum requires a previous conditional instruction which accesses the second 4KB memory region (=where the watchpoint is set), which misses in the μ TLB, and which is condition failed. The erratum also requires that no other μ TLB miss occurs between this conditional failed instruction and the unaligned access, which implies that the unaligned access must hit in the μ TLB for the first part of its access

Implications

A watchpoint does not trigger when it should.

Workaround

The erratum might occur in the case when a watchpoint is set on any of the first 3 bytes of a 4KB memory region, and unaligned accesses are not being faulted.

The workaround is then to set a guard watchpoint on the last byte of the previous page, and to deal with any false positive matches if they occur.

2.3.7 (754322) Faulty MMU translations following ASID switch

Status

Affects: Product Cortex-A9, Cortex-A9 MPCore.

Fault Type: Programmer Category B

Fault Status: Present in: All r2, r3 and r4 revisions. Open

Description

A μ TLB entry might be corrupted following an ASID switch, possibly corrupting subsequent MMU translations.

The erratum requires execution of a speculative explicit memory access, which speculation fails. The speculation fails if the memory access occurs under a mispredicted branch, or if it is conditional and the condition fails.

This speculative memory access might miss in the TLB, and cause a page table walk. The erratum occurs when the page table walk starts before the ASID switch code sequence, but completes after the ASID switch code sequence.

In this case, the μ TLB gets a new entry allocated with this new TLB entry, corresponding to the old ASID. The erratum arises from failure of the μ TLB to register the ASID value, so that MMU translations, which should happen with the new ASID following the ASID switch, might hit in this stale μ TLB entry and become corrupted.

Note that there is no Trustzone Security risk because the Security state of the access is registered in the μ TLB, and cannot be corrupted.

Implications

The erratum might cause MMU translation corruptions.

Workaround

The workaround for this erratum is to add a DSB in the ASID switch code sequence.

The ARM architecture only mandates ISB before and after the ASID switch. Adding a DSB before the ASID switch ensures that the page table walk completes before the ASID change, so that no stale entry can be allocated in the μ TLB.

Modify the examples in the *ARM Architecture Reference Manual* for synchronizing the change in the ASID and TTBR as follows:

1. The sequence:
Change ASID to 0
ISB
Change Translation Table Base Register
ISB
Change ASID to new value
Becomes:
DSB
Change ASID to 0
ISB
Change Translation Table Base Register
ISB
DSB
Change ASID to new value
2. The sequence:
Change Translation Table Base Register to the global-only mappings
ISB
Change ASID to new value
ISB

Change Translation Table Base Register to new value

Becomes:

Change Translation Table Base Register to the global-only mappings

ISB

DSB

Change ASID to new value

ISB

Change Translation Table Base Register to new value

3. And the sequence:

Set TTBCR.PD0 = 1

ISB

Change ASID to new value

Change Translation Table Base Register to new value

ISB

Set TTBCR.PD0 = 0

Becomes:

Set TTBCR.PD0 = 1

ISB

DSB

Change ASID to new value

Change Translation Table Base Register to new value

ISB

Set TTBCR.PD0 = 0

2.3.8 (764369) Data or unified cache line maintenance by MVA fails on Inner Shareable memory

Status

Affects: Product Cortex-A9 MPCore

Fault Type: Programmer Category B

Fault Status: Present in: All revisions. Open

Description

Under certain timing circumstances, a data or unified cache line maintenance operation by MVA that targets an Inner Shareable memory region might fail to propagate to either the Point of Coherency or to the Point of Unification of the system.

As a consequence, the visibility of the updated data might not be guaranteed to either the instruction side, in the case of self-modifying code, or to an external non-coherent agent, such as a DMA engine.

Conditions

The erratum requires a Cortex-A9 MPCore configuration with two or more processors, working in SMP mode, with the broadcasting of CP15 maintenance operations enabled.

The following scenario shows how the erratum can occur:

1. One CPU performs a data or unified cache line maintenance operation by MVA targeting a memory region which is locally dirty.
2. A second CPU issues a memory request targeting this same memory location within the same time frame.

A race condition can occur, resulting in the cache operation not being performed to the specified Point of Unification or Point of Coherence.

The erratum affects the following maintenance operations:

- DCIMVAC: Invalidate data or unified cache line by MVA to PoC
- DCCMVAC: Clean data or unified cache line by MVA to PoC
- DCCMVAU: Clean data or unified cache line by MVA to PoU
- DCCIMVAC: Clean and invalidate data or unified cache line by MVA to PoC.

The erratum can occur when the second CPU performs any of the following operations:

- A read request resulting from any Load instruction; the Load might be a speculative one
- A write request resulting from any Store instruction
- A data prefetch resulting from a PLD instruction; the PLD might be a speculative one.

Implications

Because it is uncertain whether execution of the cache maintenance operation propagates to either the Point of Unification or the Point of Coherence, stale data might remain in the data cache and not become visible to other agents that should have gained visibility on it.

Note that the data remains coherent on the L1 Data side. Any data read from another processor in the Cortex A9 MPCore cluster, or from the ACP, would see the correct data. In the same way, any write on the same cache line from another processor in the Cortex-A9 MPCore cluster, or from the ACP, does not cause a data corruption resulting from a loss of either data.

Consequently, the failure can only impact non-coherent agents in the systems. This can be either the instruction cache of the processor, in the case of self-modifying code, or any non-coherent external agent in the system like a DMA.

Workaround

Two workarounds are available for this erratum.

The first workaround requires the three following elements to be applied altogether:

1. Set bit[0] in the undocumented SCU Diagnostic Control register located at offset 0x30 from the PERIPHBASE address.
Setting this bit disables the *migratory bit* feature. This forces a dirty cache line to be evicted to the lower memory subsystem, which is both the Point of Coherency and the Point of Unification, when it is being read by another processor.
2. Insert a DSB instruction before the cache maintenance operation.
Note that, if the cache maintenance operation executes within a loop that performs no other memory operations, ARM recommends only adding a DSB before entering the loop.
3. Ensure there is no false sharing (on a cache line size alignment) for self-modifying code or for data produced for external non-coherent agent such as a DMA engine.
For systems which cannot prevent false sharing in these regions, this third step can be replaced by performing the sequence of DSB followed by Cache maintenance operation twice.

Note that even when all three components of the workaround are in place, the erratum might still occur. However, this would require some extremely rare and complex timing conditions, so that the probability of reaching the point of failure is extremely low. This, and the fact that the erratum requires an uncommon software scenario, explains why this workaround is likely to be a reliable practical solution for most systems.

To ARM's knowledge, no failure has been observed in any system when all three components of this workaround have been implemented.

For critical systems that cannot cope with the extremely low failure risks associated with the above workaround, a second workaround is possible which involves changing the mapping of the data being accessed so that is in a Non-Cacheable area. This ensures that the written data remains uncached. This means it is always visible to non-coherent agents in the system, or to the instruction side in the case of self-modifying code, without any need for cache maintenance operation.

2.3.9 (775420) A data cache maintenance operation which aborts, followed by an ISB, without any DSB in-between, might lead to deadlock

Status

Affects: product Cortex-A9, Cortex-A9 MPCore.

Fault Type: Programmer Category B

Fault Status: Present in: r2p2, r2p3, r2p4, r2p6, r2p8, r2p10, r3p0. Fixed in r4p0

Description

Under certain micro-architectural circumstances, a data cache maintenance operation which aborts, followed by an ISB, with no DSB occurring between these events, might lead to processor deadlock.

Conditions

The erratum requires the following conditions:

1. Some write operations are being handled by the processor, and take a long time to complete. The typical situation is when the write operation (STR, STM, ...) has missed in the L1 Data Cache.
2. No memory barrier (DMB or DSB) is inserted between the write operation and the data cache maintenance operation mentioned in condition 3.
3. A data cache maintenance operation is performed, which aborts because of its MMU settings.
4. No memory barrier (DMB or DSB) is inserted between the data cache maintenance operation in condition 3 and the ISB in condition 5. Any other kind of code can be executed here, starting with the abort exception handler following the aborted cache maintenance operation.
5. An ISB instruction is being executed by the processor.
6. No memory barrier (DMB or DSB) is inserted between the ISB in condition 5 and the read or write operation in condition 7.
7. A read or write operation is executed.

With the above conditions, an internal Data Side drain request signal might remain sticky, causing the ISB to wait for the Data Side to be empty, which never happens because the last read or write operation waits for the ISB to complete.

Implications

The erratum can lead to processor deadlock.

Workaround

A simple workaround for this erratum is to add a DSB at the beginning of the abort exception handler.

2.3.10 (782772) Speculative execution of a Load-Exclusive or Store-Exclusive instruction after a write to Strongly Ordered memory might deadlock the processor

Status

Affects: product Cortex-A9, Cortex-A9 MPCore.

Fault Type: Programmer Category B

Fault Status: Present in: All r1, r2 and r3 revisions. Fixed in r4p0

Description

Under certain timing circumstances, a processor might deadlock when the execution of a write to a Strongly Ordered memory region is followed by the speculative execution of a Load-Exclusive or a Store-Exclusive instruction that is mis-speculated.

The mis-speculation can be due to either the Load-Exclusive or Store-Exclusive instruction being conditional, and failing its condition code check, or to the Load-Exclusive or Store-Exclusive instruction being speculatively executed in the shadow of a mispredicted branch.

Configurations affected

This erratum affects all configurations of the processor.

Conditions

The erratum requires the following conditions:

- The processor executes a write instruction to a Strongly Ordered memory region.
- The processor speculatively executes a Load-Exclusive or Store-Exclusive instruction that is either:
 1. A conditional instruction
 2. An instruction in the shadow of a conditional branch.
- The Load-Exclusive or Store-Exclusive instruction is cancelled because the speculation was incorrect, because either:
 1. The conditional Load-Exclusive or Store-Exclusive instruction failed its condition-code check
 2. The conditional branch was mispredicted, so that all subsequent instructions speculatively executed must be flushed, including the Load-Exclusive or Store-Exclusive.

The erratum also requires additional timing conditions to be met. These are specific to each platform, and are not controllable by software. These timing conditions includes the fact that the response to the Strongly Ordered write from the external memory system must be received at the same time as the mis-speculation is identified in the processor.

Implications

The erratum causes processor deadlock.

Workaround

The recommended workaround is to place a DMB instruction before each Load-Exclusive / Store-Exclusive loop sequence, to ensure that no pending write request can interfere with the execution of the Load-Exclusive or Store-Exclusive instructions. The implementation of this workaround can be restricted to code regions which have access to Strongly Ordered memory.

2.3.11 (782773) Updating a translation entry to move a page mapping might erroneously cause an unexpected translation fault

Status

Affects: product Cortex-A9, Cortex-A9 MPCore.

Fault Type: Programmer Category B

Fault Status: Present in: All r0, r1, r2 and r3 revisions. Fixed in r4p0

Description

Under certain conditions specific to the Cortex-A9 micro-architecture, a write operation that updates a Cacheable translation table entry might cause both the old and the new translation entry to be temporarily invisible to translation table walks, thus erroneously causing a translation fault.

Conditions

The erratum requires the following conditions to happen:

1. The processor has its Data Cache and MMU enabled.
2. The TTB registers are set to work on Cacheable descriptors memory regions.
3. The processor is updating an existing Cacheable translation table entry, and this write operation hits in the L1 Data Cache.
4. A hardware translation table walk is attempted. The hardware translation table walk can be due to either an Instruction fetch, or to any other instruction execution that requires an address translation, including any load or store operation. This hardware translation walk must attempt to access the entry being updated in condition 2, and that access must hit in the L1 Data Cache.

In practice, this scenario can happen when an OS is changing the mapping of a physical page. The OS might have an existing mapping to a physical page (the old mapping), but wants to move the mapping to a new page (the new mapping). To do this, the OS might:

1. Write a new translation entry, without cancelling the old one. At this point the physical page is accessible using either the old mapping or the new mapping.
2. Execute a DSB instruction followed by an ISB instruction pair, to ensure that the new translation entry is fully visible.
3. Remove the old entry.

Because of the erratum, this sequence might fail because it can happen that neither the new mapping, nor the old mapping, is visible after the new entry is written, causing a Translation fault.

Implications

The erratum causes a Translation fault.

Workaround

The recommended workaround is to perform a clean and invalidate operation on the cache line that contains the translation entry before updating the entry, to ensure that the write operation misses in the Data Cache. This workaround prevents the micro-architectural conditions for the erratum from happening. Interrupts must be temporarily disabled so that no interrupt can be taken between the maintenance operation and the translation entry update. This avoids the possibility of the interrupt service routine bringing the cache line back in the cache.

Another possible workaround is to place the translation table entries in Non-Cacheable memory areas, but this workaround is likely to have a noticeable performance penalty.

Note that inserting a DSB instruction immediately after writing the new translation table entry significantly reduces the probability of hitting the erratum, but is not a complete workaround.

2.3.12 (794072) A short loop including a DMB instruction might cause a denial of service on another processor which executes a CP15 broadcast operation

Status

Affects: Cortex-A9 MPCore.

Fault Type: Programmer Category B

Fault Status: Present in: All r1, r2, r3 and r4 revisionsOpen

Description

A processor which continuously executes a short loop containing a DMB instruction might prevent a CP15 operation broadcast by another processor making further progress, causing a denial of service.

Configurations affected

This erratum affects all Cortex-A9 MPCore processors with two or more processors.

Conditions

The erratum requires the following conditions:

- Two or more processors are working in SMP mode (ACTLR.SMP=1)
- One of the processors continuously executes a short loop containing at least one DMB instruction.
- Another processor executes a CP15 maintenance operation that is broadcast. This requires that this processor has enabled the broadcasting of CP15 operations (ACTLR.FW=1)

For the erratum to occur, the short loop containing the DMB instruction must meet all of the following additional conditions:

- No more than 10 instructions other than the DMB are executed between each DMB
- No non-conditional Load or Store, or conditional Load or Store which pass the condition code check, are executed between each DMB

When all the conditions for the erratum are met, the short loop is creating a continuous stream of DMB instructions. This might cause a denial of service, by preventing the processor executing the short loop from executing the received broadcast CP15 operation. As a result, the processor that originally executed the broadcast CP15 operation is stalled until the execution of the loop is interrupted.

Note that because the process issuing the CP15 broadcast operation cannot complete operation, it cannot enter any debug-mode, and cannot take any interrupt. If the processor executing the short loop also cannot be interrupted, for example if it has disabled its interrupts, or if no interrupts are routed to this processor, this erratum might cause a system livelock.

Implications

The erratum might create performance issues, or in the worst case it might cause a system livelock if the processor executing the DMB is in an infinite loop that cannot be interrupted.

Workaround

This erratum can be worked round by setting bit[4] of the undocumented Diagnostic Control Register to 1. This register is encoded as CP15 c15 0 c0 1.

This bit can be written in Secure state only, with the following Read/Modify/Write code sequence:

```
MRC p15,0,rt,c15,c0,1
```

```
ORR rt,rt,#0x10  
MCR p15,0,rt,c15,c0,1
```

When it is set, this bit causes the DMB instruction to be decoded and executed like a DSB.

Using this software workaround is not expected to have any impact on the overall performance of the processor on a typical code base.

Other workarounds are also available for this erratum, to either prevent or interrupt the continuous stream of DMB instructions that causes the deadlock. For example:

- Inserting a non-conditional Load or Store instruction in the loop between each DMB
- Inserting additional instructions in the loop, such as NOPs, to avoid the processor seeing back to back DMB instructions.
- Making the processor executing the short loop take regular interrupts.

2.3.13 (794073) Speculative instruction fetches with MMU disabled might not comply with architectural requirements

Status

Affects: product Cortex-A9, Cortex-A9 MPCore.

Fault Type: Programmer Category B

Fault Status: Present in: All revisions. Open

Description

When the MMU is disabled, an ARMv7 processor must follow some architectural rules regarding speculative fetches and the addresses to which these can be initiated. These rules avoid potential read accesses to read-sensitive areas. For more information about these rules see the description of "Behavior of instruction fetches when all associated MMUs are disabled" in the *ARM Architecture Reference Manual*, ARMv7-A and ARMv7-R edition.

A Cortex-A9 processor usually operates with both the MMU and branch prediction enabled. If the processor operates in this condition for any significant amount of time, the *Branch Target Address Cache* (BTAC) will contain branch predictions. If the MMU is then disabled, but branch prediction remains enabled, these stale BTAC entries can cause the processor to violate the rules for speculative fetches.

Configurations affected

This erratum affects all configurations of the processor.

Conditions

The erratum can occur only if the following sequence of conditions is met:

1. MMU and branch prediction are enabled.
2. Branches are executed.
3. MMU is disabled, and branch prediction remains enabled.

Implications

If the above conditions occur, it is possible that after the MMU is disabled, speculative instruction fetches might occur to read-sensitive locations.

Workaround

The recommended workaround is to invalidate all entries in the BTAC, by executing an *Invalidate Entire Branch Prediction Array* (BPIALL) operation followed by a DSB, before disabling the MMU.

Another possible workaround is to disable branch prediction when disabling the MMU, and keep branch prediction disabled until the MMU is re-enabled.

2.3.14 (794074) A write request to Uncacheable, Shareable normal memory region might be executed twice, possibly causing a software synchronisation issue

Status

Affects: product Cortex-A9, Cortex-A9 MPCore.

Fault Type: Programmer Category B

Fault Status: Present in: All revisions. Open

Description

Under certain timing circumstances specific to the Cortex-A9 microarchitecture, a write request to an Uncacheable, Shareable Normal memory region might be executed twice, causing the write request to be sent twice on the AXI bus. This might happen when the write request is followed by another write into the same naturally aligned doubleword memory region, without a DMB between the two writes.

The repetition of the write usually has no impact on the overall behaviour of the system, unless the repeated write is used for synchronisation purposes.

Configurations affected

The erratum affects all configurations of the processor.

Conditions

The erratum requires the following conditions:

1. A write request is performed to an Uncacheable, Shareable Normal memory region.
2. Another write request is performed into the same naturally doubleword aligned memory region. This second write request must not be performed to the exact same bytes as the first store.

A write request to Normal memory region is treated as Uncacheable in the following cases:

- The write request occurs while the Data Cache is disabled.
- The write request is targeting a memory region marked as Normal Memory Non-Cacheable or Cacheable Write-Through.
- The write request is targeting a memory region marked as Normal Memory Cacheable Write-Back and Shareable, and the CPU is in AMP mode.

Implications

This erratum might have implications in a multi-master system where control information is passed between several processing elements in memory using a communication variable, for example a semaphore. In such a system, it is common for communication variables to be claimed using a Load-Exclusive/Store-Exclusive, but for the communication variable to be cleared using a non-Exclusive store. This erratum means that the clearing of such a communication variable might occur twice. This might lead to two masters apparently claiming a communication variable, and therefore might cause data corruption to shared data.

A scenario in which this might happen is:

```
MOV r1,#0x40; address is double-word aligned, mapped in
                ; Normal Non-cacheable Shareable memory
Loop:LDREXr5, [r1,#0x0]; read the communication variable
CMP r5, #0 ; check if 0
STREXEQ r5, r0, [r1]; attempt to store new value
CMPEQ r5, #0; test if store succeeded
BNE Loop; retry if not
DMB ; ensures that all subsequent accesses are observed when
```

```
; gaining of the communication variable has been observed
; loads and stores in the critical region can now be performed
MOV r2, #0
MOV r0, #0
DMB          ; ensure all previous accesses are observed before the
; communication variable is cleared
STR r0, [r1]; clear the communication variable with normal store
STR r2, [r1, #0x4]
; previous STR might merge and be sent again, which might
; cause undesired release of the communication variable.
```

This scenario is valid when the communication variable is a byte, a half-word, or a word

Workaround

There are several possible workarounds:

- Add a DMB after clearing a communication variable:
STR r0, [r1]; clear the communication variable
DMB ; ensure the previous STR is complete
Also any IRQ or FIQ handler must execute a DMB at the start to ensure as well the clear of any communication variable is complete.
- Ensure there is no other data using the same naturally aligned 64-bit memory location as the communication variable:
ALIGN 64
communication_variable DCD 0
unused_data DCD 0
LDR r1, = communication_variable
- Use a Store-Exclusive to clear the communication variable, rather than a non-Exclusive store.

2.4 Category B (Rare)

This category includes the following errata:

2.4.1 (743622) Faulty logic in the Store Buffer might cause data corruption

Status

Affects: Product Cortex-A9, Cortex-A9 MPCore.

Fault Type: Programmer Category B (Rare)

Fault Status: Present in: All r2 revisions Fixed in r3p0

Description

Under very rare conditions, a faulty optimization in the Cortex-A9 Store Buffer might cause data corruption.

Conditions

The code sequence which exhibits the failure requires at least five Cacheable writes in 64-bit data chunk:

- three of the writes must be in the same cache line
- another write must be in a different cache line.

All of the above four writes hit in the L1 data cache.

A fifth write which fully writes a 64-bit data chunk is required in either of the above two cache lines.

With the above code sequence, under very rare circumstances, this fifth write might be corrupted, with the written data either being lost, or being written in another cache line.

The conditions under which the erratum can occur are extremely rare, and require the coincidence of multiple events and states in the Cortex-A9 micro-architecture.

For example, assume A, A', A'', and A''' are all in the same cache line, and that B and B' are in another cache line. The following code sequence might trigger the erratum:

- STR A
- STR A'
- STR A''
- STR B
- STR A''' (or STR B')

At the time when the first four STRs are in the Cortex-A9 Store Buffer, and the fifth STR arrives at a very precise cycle in the Store Buffer input stage, then the fifth STR might not see its cache line dependency on the previous STR instructions. Because of this, in cases when the cache line A or B gets invalidated because of a coherent request from another CPU, the fifth STR might write in a faulty cache line, which causes data corruption.

An alternative version of the erratum might happen even without a coherent request. This is the case when the fifth STR is a 64-bit write in the same location as one of A, A', A''. In these circumstances the erratum might also occur. Note that this is quite an uncommon scenario because it requires a first write to a memory location which is then immediately and fully overwritten.

Implications

When this erratum occurs, it causes data corruption.

Workaround

A software workaround is to set bit[6] in the undocumented Diagnostic Control register placed in CP15 c15 0 c0 1.

You can write this bit in Secure state only, by using the following Read/Modify/Write code sequence:

```
MRC p15,0,rt,c15,c0,1
ORR rt,rt,#0x40
MCR p15,0,rt,c15,c0,1
```

The 'fast lookup' optimization in the Store Buffer is disabled when this bit is set. This prevents the erratum occurring.

Setting this bit has no visible impact on the overall performance or power consumption of the processor.

2.4.2 (751473) Under very rare circumstances, automatic data prefetcher might cause deadlock or data corruption

Status

Affects: Product Cortex-A9, Cortex-A9 MPCore.

Fault Type: Programmer Category B (Rare)

Fault Status: Present in: All r0, r1 and r2 revisions Fixed in r3p0

Description

Under very rare timing circumstances, the automatic data prefetcher might cause address hazard issues, possibly leading to a data corruption or a processor deadlock.

Conditions

This erratum can occur only when the Data Cache and MMU are enabled in the following cases:

- on all memory regions marked as Write-Back Non-Shared, when the data prefetcher in L1 is enabled (ACTLR[2]=1'b1), regardless of the ACTLR.SMP bit
- on all memory regions marked as Write-Back Shared when the data prefetch hint in L2 is enabled (ACTLR[1]=1'b1) and when the processor is in SMP mode (ACTLR.SMP=1'b1).

Implications

When this erratum occurs, a data corruption or a processor deadlock can occur.

Workaround

The workaround for this erratum is to not enable the automatic data prefetcher by keeping ACTRL[2:1]=2'b00, which is the default value on exit from reset.

Although this feature might show significant performance gain on a few synthetic benchmarks, it usually has no impact in real systems. Therefore, this workaround is unlikely to cause any visible impact on final products.

2.4.3 (761320) Full cache line writes to the same memory region from at least two processors might deadlock the processor

Status

Affects: Product Cortex-A9 MPCore.

Fault Type: Programmer Category B (Rare)

Fault Status: Present in: All r0, r1, r2 and r3 revisions. Fixed in r4p0

Description

Under very rare circumstances, full cache line writes from at least two processors on cache lines in hazard with other cache accesses might cause arbitration issues in the SCU, leading to processor deadlock.

Configurations affected

This erratum affects the configurations of the processor with three or more active coherent agents, which is either:

- Two or more processors if the ACP is present.
- Three or more processors.

Conditions

To trigger the erratum, at least three agents need to be working in SMP mode, and accessing coherent memory regions.

Two or more processors need to perform full cache line writes, to cache lines which are in hazard with other access requests in the SCU. The hazard in the SCU happens when another processor, or the ACP, is performing a read of or a write to the same cache line.

The following example describes one scenario that might cause this deadlock:

- CPU0 performs a full cache line write to address A, then a full cache line write to address B
- CPU1 performs a full cache line write to address B, then a full cache line write to address A
- CPU2 performs read accesses to addresses A and B

Under certain rare timing circumstances, the requests might create a loop of dependencies, causing a processor deadlock.

Implications

When the erratum happens, it leads to system deadlock.

It is important to note that any scenario leading to this deadlock situation is uncommon. It requires two processors writing full cache lines to a coherent memory region, without taking any semaphore, with another processor or the ACP accessing the same lines at the same time, meaning that these latter accesses are not deterministic. This, combined with the extremely rare microarchitectural timing conditions under which the defect can happen, explains why the erratum is not expected to cause any significant malfunction in real systems.

Workaround

This erratum can be worked round by setting bit[21] of the undocumented Diagnostic Control Register to 1. This register is encoded as CP15 c15 0 c0 1.

The bit can be written in Secure state only, with the following Read/Modify/Write code sequence:

MRC p15,0,rt,c15,c0,1

```
ORR rt,rt,#0x200000  
MCR p15,0,rt,c15,c0,1
```

When this bit is set, the "direct eviction" optimization in the Bus Interface Unit is disabled, which means this erratum cannot occur.

Setting this bit might prevent the Cortex-A9 from utilizing the full bandwidth when performing intensive full cache line writes, and therefore a slight performance drop might be visible.

In addition, this erratum cannot occur if at least one of the following bits in the Diagnostic Control Register is set to 1:

- bit [23] - Disable Read-Allocate mode
- bit [22] - Disable Write Allocate Wait mode

2.5 Category C

This section describes Category C errata.

2.5.1 (725631) ISB is counted in Performance Monitor events 0x0C and 0x0D

Status

Affects: Product Cortex-A9, Cortex-A9 MPCore.

Fault Type: Programmer Category C

Fault Status: Present in: All revisions. Open

Description

The ISB is implemented as a branch in the Cortex-A9 micro-architecture.

This implies that events 0x0C (software change of PC) and 0x0D (immediate branch) are asserted when an ISB occurs, which is not compliant with the ARM Architecture.

Implications

The count of events 0x0C and 0x0D are not completely precise when using the Performance Monitor counters, because the ISB is counted together with the real software changes to PC (for 0x0C) and immediate branches (0x0D).

The erratum also causes the corresponding PMUEVENT bits to toggle in case an ISB executes.

- PMUEVENT[13] relates to event 0x0C
- PMUEVENT[14] relates to event 0x0D.

Workaround

You can count ISB instructions alone with event 0x90.

You can subtract this ISB count from the results you obtained in events 0x0C and 0x0D, to obtain the precise count of software change of PC (0x0C) and immediate branches (0x0D).

2.5.2 (729808) PLE Stride ranges from 0 to 255, instead of from 1 to 256

Status

Affects: Product Cortex-A9, Cortex-A9 MPCore.

Fault Type: Programmer Category C

Fault Status: Present in: r2p0, r2p1, r2p5, r2p7, r2p9. Fixed in all other revisions.

Description

The PLE stride is an 8-bit bitfield that is specified in bits [17:10] when programming a new PLE channel.

This PLE stride is expected to range from 8'b00000000, to prefetch contiguous blocks, up to 8'b11111111, to prefetch blocks every 256 words.

Because of the erratum, the PLE stride is the expected PLE stride minus 1. That is, a PLE stride of 8'b00000000 means a real stride of 0, instead of the expected stride of 1 word to prefetch contiguous blocks. In a similar way, a PLE stride of 8'b11111111 prefetches blocks every 255 words, instead of the expected 256 words.

Implications

Because of the erratum, it is possible to program the PLE stride only from 0 to 255, instead of from 1 to 256 as expected.

Workaround

The software workaround is to program the PLE stride with +1, compared with the original value that would have been programmed without the erratum.

2.5.3 (729814) CFGSDISABLE pin value modification might not be detected

Status

Affects: Product Cortex-A9 MPCore.

Fault Type: Programmer Category C

Fault Status: Present in: r2p0, r2p1, r2p5, r2p7, r2p9. Fixed in all other revisions.

Description

The CFGSDISABLE input pin is sampled internally in the Cortex-A9 MPCore in the IC (Interrupt Controller) module. When the IC clock-gating is enabled, the clock on this sampling register might be disabled, meaning that the CFGSDISABLE input pin is not sampled correctly. This does not create any issue if the CFGSDISABLE input value does not change, but if it does change, the old value of CFGSDISABLE might be used instead of the new one.

Conditions

- IC dynamic clock-gating is enabled (SCU.Control Register bit[6], IC standby enable, is set to 1)
- CFGSDISABLE input value is toggling.

Implications

The CFGSDISABLE input is used to disable the access to some security-critical configuration registers. Because of the erratum, this value might be faulty, creating a potential security hole.

Workaround

Two software workarounds are available for this erratum:

- Make sure the SCU Control register bit[6] remains at 0, which is the default (reset) value. This ensures the IC dynamic clock gating is never enabled.
- Perform a read of the Distributor Enable register after the CFGSDISABLE pin toggles. Doing the read requires the clock to be enabled on the register, thus sampling the correct CFGSDISABLE value.

2.5.4 (729816) TLB Lockdown write entry operation might not execute

Status

Affects: Product Cortex-A9, Cortex-A9 MPCore.

Fault Type: Programmer Category C

Fault Status: Present in: r2p0, r2p1, r2p5, r2p7, r2p9. Fixed in all other revisions.

Description

Under certain timing conditions specific to the Cortex-A9 micro-architecture, a CP15 write operation into the Lockdown TLB entries might not execute, leaving the TLB Lockdown entry in its previous state.

Implications

Because the TLB entry is not written, any future code that relies on this entry might be faulty.

Workaround

The workaround for this erratum is to perform a TLB Lockdown Read operation after a TLB Lockdown Write operation, to ensure that the TLB Write operation is successful.

If it is not, the code should repeat the TLB Lockdown write operation until the register write is successful.

2.5.5 (729817) Main ID register alias addresses are not mapped on Debug APB interface

Status

Affects: Product Cortex-A9, Cortex-A9 MPCore.

Fault Type: Programmer Category C

Fault Status: Present in: All revisions. Open

Description

The ARM Debug Architecture specifies registers 838 and 839 as *Alias of the Main ID register*. They should be accessible using the APB Debug interface at addresses 0xD18 and 0xD1C.

The two alias addresses are not implemented in Cortex-A9. A read access at either of these two addresses returns 0, instead of the MIDR value.

Note

Read accesses to these two registers using the internal CP14 interface are trapped to UNDEFINED, which is compliant with the ARM Debug architecture. Therefore the erratum only applies to the alias addresses using the external Debug APB interface.

Implications

If the debugger, or any other external agent, tries to read the MIDR register using the alias addresses, it will get a faulty answer (0x0), which can cause indeterminate errors in the debugger afterwards.

Workaround

The workaround for this erratum is to always access the MIDR at its original address, 0xD00, and not to use its alias address.

2.5.6 (729818) In debug state, next instruction is stalled when sdabort flag is set, instead of being discarded

Status

Affects: Product Cortex-A9, Cortex-A9 MPCore.

Fault Type: Programmer Category C

Fault Status: Present in: All revisions. Open.

Description

When the processor is in debug state, an instruction written to the ITR after a Load/Store instruction that aborts gets executed on clearing the SDABORT_1, instead of being discarded.

Conditions

- Debugger has put the extDCCmode bits into Stall mode
- A previously issued load/store instruction has generated a synchronous Data Abort (for example, an MMU fault)
- For efficiency, the debugger does not read DBGDSCRExt immediately, to see if the load/store has completed and has not aborted, but writes further instructions to the ITR, expecting them to be discarded if a problem occurs
- The debugger reads the DBGDSCR at the end of the sequence and discovers the load/store aborted
- The debugger clears the SDABORT_1 flag (by writing to the Clear Sticky Aborts bit in DBGDRCR).

Under these conditions, the instruction that follows in the ITR might execute instead of being discarded.

Implications

Indeterminate failures can occur because of the instruction being executed when it should not. In most cases, it is unlikely that the failure will cause any significant issue.

Workaround

There are a selection of workarounds with increasing complexity and decreasing impact. In each case the impact is a loss of performance when debugging:

1. Do not use stall mode.
2. Do not use stall mode when doing load/store operations.
3. Always check for a sticky abort after issuing a load/store operation in stall mode (the cost of this probably means workaround number 2 is a preferred alternative).
4. Always check for a sticky abort after issuing a load/store operation in stall mode before issuing any further instructions that might corrupt an important target state (such as further load/store instructions, instructions that write to live registers such as VFP, CP15).

2.5.7 (730345) PLE Wait States resolution is incorrect

Status

Affects: Product Cortex-A9, Cortex-A9 MPCore.

Fault Type Programmer Category C

Fault Status: Present in: r2p0, r2p1, r2p5, r2p7, r2p9. Fixed in all other revisions.

Description

PLEPCR[7:0] bitfield contains the PLE Wait States parameter, which permits the Privilege modes to limit the issuing rate of PLD requests performed by the PLE engine to prevent saturation of the external memory bandwidth.

PLEPCR[7:0] is supposed to be the real number of wait states, to provide a resolution of 1 cycle.

The erratum is that the number of wait states used by the PLE between two PLD requests is PLEPCR[7:0] shifted left by 4 bits, which provides a resolution of 16 cycles. That is, the number of wait states is $(16 * \text{PLEPCR}[7:0] + 1)$, instead of $(\text{PLEPCR}[7:0] + 1)$.

Implications

Because of the erratum, PLE Wait States between two PLD requests has a resolution of 16 cycles instead of 1. Although this is not likely to create any significant issue, any software developer should be aware of this erratum to fine tune their PLE programming.

Workaround

To workaround this erratum, the programmer needs to compensate for the hardware “shift left by 4” of the PLEPCR[7:0] value. That is, it is necessary to program PLEPCR[7:0] with a value shifted right by 4 compared with the original value which would have been programmed without the erratum.

————— **Note** —————

This solution provides a PLE resolution programming of 16 cycles, and there is no workaround to narrow this down.

2.5.8 (740661) Event 0x74 / PMUEVENT[38:37] might be inaccurate

Status

Affects: Product Cortex-A9, Cortex-A9 MPCore

Fault Type: Programmer Category C

Fault Status: Present in: All r1 and r2 revisions. Fixed in r3p0

Description

Event 0x74 counts the total number of Neon instructions passing through the Register Rename pipeline stage. Because of the erratum, the stall information is not taken into account. Therefore, one Neon instruction that remains n cycles in the Register Rename stage is counted as n Neon instructions.

As a consequence, the count of event 0x74 might be corrupted, and cannot be relied upon.

The event is also reported externally on PMUEVENT[38:37], which suffers from the same inaccuracy.

Implications

The implication of this erratum is that Neon instructions cannot be counted reliably in the versions of the product which are affected by this erratum.

Workaround

No workaround is possible to achieve the required functionality of counting how many Neon instructions are executed (or renamed) in the processor.

2.5.9 (740663) Event 0x68 / PMUEVENT[9:8] might be inaccurate

Status

Affects: Product Cortex-A9, Cortex-A9 MPCore.

Fault Type: Programmer Category C

Fault Status: Present in: All r0, r1 and r2 revisions. Fixed in r3p0

Description

Event 0x68 counts the total number of instructions passing through the Register Rename pipeline stage.

Under certain conditions, some branch-related instructions might pass through this pipeline stage without being counted. As a consequence, event 0x68 might be inaccurate, with a lower value than expected.

The event is also reported externally on PMUEVENT[9:8], which suffers from the same inaccuracy.

Conditions

- Events are enabled.
- One of the PMU counters is programmed to count event 0x68 - number of instructions passing through the Register Rename stage. Alternatively, an external component counts, or relies on, PMUEVENT[9:8].
- A program executes, which contains one of the following instruction:
 - A Branch immediate, without Link
 - An ISB instruction
 - An HB instruction, without Link and without parameter, in ThumbEE state
 - An ENTERX or LEAVEX instruction, in Thumb or ThumbEE state.
- The program executed causes some stalls in the processor pipeline.

Under certain timing conditions, specific to the Cortex-A9 micro-architecture, a cycle stall in the processor pipeline might hide the instructions mentioned above, resulting in a corrupted count for event 0x68, or a corrupted value on PMUEVENT[9:8] during this given cycle. If the hidden instruction appears in a loop, the count difference can be significant.

As an example, consider the following loop:

```
loop    subs            r0, #1
        vadd.f32 q1, q1, q0
        bne            loop
```

The loop contains three instructions, so the final instruction count should (approximately) be three times the number of executed loops. In practice, the vadd causes a pipeline stall after a few loop iterations. This stall hides the branch instruction (bne loop), so that only two instructions are counted per loop, and the final count appears as twice the number of executed loops instead of three times this number.

Implications

The implication of this erratum is that the values of event 0x68 and PMUEVENT[9:8] are imprecise, and cannot be relied upon.

Workaround

No workaround is possible to achieve the required functionality of counting precisely how many instructions are passing through the Register Rename pipeline stage.

2.5.10 (743623) Bad interaction between a minimum of seven PLDs and one Non-Cacheable LDM, can lead to deadlock

Status

Affects: Product Cortex-A9, Cortex-A9 MPCore

Fault Type: Programmer Category C

Fault Status: Present in: All r0, r1 and r2 revisions. Fixed in r3p0

Description

Under very rare circumstances, a deadlock can occur in the processor when handles a minimum of seven PLD instructions, shortly followed by one LDM to an uncachable memory location.

The LDM is treated as Uncacheable in the following cases:

1. The LDM occurs while the Data Cache is Off.
2. The LDM targets a memory region marked as Strongly Ordered, Device, Normal Memory Non-Cacheable, or Normal Memory Write-Through
3. The LDM targets a memory region marked as Shareable Normal Memory Write-Back, and the CPU is in AMP mode.

Conditions

The code sequence that exhibits this erratum requires at least seven PLDs shortly followed by one LDM to an uncachable memory region. The erratum occurs when the LDM appears on the AXI bus before any of the seven PLDs, which can possibly happen when the first PLD misses in the μ TLB. In this case it needs to perform a TLB request, which may not be serviced immediately because the main TLB is already performing a Page Table Walk for another resource (for example, Instruction side), or because the PLD request to the main TLB is missing and causing a Page Table Walk.

Also note that the conditions above are not sufficient to recreate the failure, because additional rare conditions on the internal state of the processor are necessary to trigger the errata.

Implications

The erratum might create a processor deadlock. However, the conditions which are required for this to occur are extremely unlikely to occur in real code sequences.

Workaround

The primary workaround is to avoid using the offending code sequence, that is, do not use uncachable LDM at the same time as an intensive use of PLD instructions.

If not possible, another workaround for this erratum is to set bit[20] in the undocumented Control register, which is placed in CP15 c15 0 c0 1.

This bit needs to be written with the following Read/Modify/Write code sequence:

```
MRC      p15,0,r0,c15,c0,1
ORR      r0,r0,#0x00100000
MCR      p15,0,r0,c15,c0,1
```

Setting this bit causes all PLD instructions to be treated as NOPs, with the consequence that code sequences usually using the PLDs, such as the memcpy() routine, might suffer from a visible performance drop.

2.5.11 (743625) A coherent ACP request might interfere with a non-cacheable SWP/SWPB from the processor, potentially causing deadlock

Status

Affects: Product Cortex-A9 MPCore.

Fault Type: Programmer Category C

Fault Status: Present in: All r0, r1 and r2 revisions. Fixed in r3p0

Description

Under very rare circumstances, a faulty address hazard checking in the SCU might cause a coherent ACP request to badly interfere with a Non-Cacheable SWP/SWPB request performed by the CPU.

Conditions

This erratum can only occur in a Cortex-A9 MPCore product containing one single processor, with the ACP present. The processor must work in coherent mode, with ACTLR.SMP=1, and with its data cache enabled.

The processor also needs to have enabled SWP instructions, by setting ACTLR.SWP=1. SWP and SWPB instructions are strongly deprecated by ARM, and in the Cortex-A9 these instructions are disabled by default.

To reproduce the failure, the following events are required:

1. The processor performs a write in a cache line A, marked as coherent (Write-Back Shared).
2. This line exits the processor, either because the processor has switched into Read-Allocate mode, so that the write is immediately performed externally, or because the cache line is naturally evicted, replaced by another line, or cleaned.
3. The processor then performs another write in this cache line A, which brings the data back in dirty state in its data cache.
4. Finally, the processor executes a SWP or SWPB instruction to a Non-Cacheable region (Strongly Ordered, Device, or Normal Memory Non-Cacheable).

The offending sequence also requires the ACP to perform a coherent request in cache line A. In r0p0, r0p1 and r0p2, the ACP request is coherent if AxUSER[0]=1. Since r1p0 revision, the ACP request is coherent if AxUSER[0]=1 and AxCACHE[1]=1.

Under these conditions, and if the CPU has not performed any other write since cache line A was written externally, then a faulty address hazard might be detected in the SCU, which prevents all current transactions completing.

Implications

Under certain timing conditions, specific to the Cortex-A9 MPCore micro-architecture, the erratum can lead to processor deadlock.

Workaround

The recommended workaround is to use LDREX/STREX instead of SWP/SWPB instructions, which are deprecated in the ARMv7 architecture.

In case SWP and SWPB cannot be suppressed from the code, an alternative workaround is to perform a dummy write to a Non-Shared, Non-Cacheable memory location before executing the SWP or SWPB, which would clear the faulty hazard checking.

2.5.12 (743626) An imprecise external abort received while the processor enters WFI may cause a processor deadlock

Status

Affects: Product Cortex-A9, Cortex-A9 MPCore.

Fault Type: Programmer Category C

Fault Status: Present in: All r0, r1 and r2 revisions. Fixed in r3p0

Description

An imprecise external abort received while the processor is ready to enter into WFI state might cause a processor deadlock.

Explicit memory transactions can be completed by inserting a DSB before the WFI instruction, but this does not prevent memory accesses generated by:

- Previously issued PLD instructions
- Page table walks associated with previously issued PLD instructions or as a result of the PLE engine.

If an external abort returns as a result of one of these memory accesses after executing a WFI instruction, the processor can deadlock.

Implications

In case the non-explicit memory request receives an external imprecise abort response while the processor is ready to enter into WFI state, the processor might deadlock.

In practical systems, it is not likely that these memory transactions will generate an external abort, because external aborts are usually a sign of significant corruption in the system.

Workaround

The workaround for this erratum is to protect all memory regions which can return an imprecise external abort with the correct MMU settings, to prevent any external aborts.

2.5.13 (751471) DBGPCSR format is incorrect

Status

Affects: Product Cortex-A9, Cortex-A9 MPCore.

Fault Type: Programmer Category C

Fault Status: Present in: All revisions. Open

Description

With respect to the DBGPCSR register, the ARM Architecture specifies that:

- DBGPCSR[31:2] contain the sampled value of bits [31:2] of the PC.
- The sampled value is an instruction address plus an offset that depends on the processor instruction set state.
- DBGPCSR[1:0] contain the meaning of PC Sample Value, with the following permitted values:
 - 0b00 ((DBGPCSR[31:2] << 2) – 8) references an ARM state instruction
 - 0bx1 ((DBGPCSR[31:1] << 1) – 4) references a Thumb or ThumbEE state instruction
 - 0b10 IMPLEMENTATION DEFINED.

This field encodes the processor instruction set state, so that the profiling tool can calculate the true instruction address by subtracting the appropriate offset from the value sampled in bits [31:2] of the register.

In Cortex-A9, the DBGPCSR samples the target address of executed branches (but possibly still speculative to data aborts), with the following encodings:

- DBGPCSR[31:2] contain the address of the target branch instruction, with no offset.
- DBGPCSR[1:0] contains the execution state of the target branch instruction:
 - 0b00 for an ARM state instruction
 - 0b01 for a Thumb state instruction
 - 0b10 for a Jazelle state instruction
 - 0b11 for a ThumbEE state instruction

Implications

The implication of this erratum is that the debugger tools must not rely on the architected description for the value of DBGPCSR[1:0], nor remove any offset from DBGPCSR[31:2], to obtain the expected PC value.

Subtracting 4 or 8 from the DBGPCSR[31:2] value would lead to an area of code which is unlikely to have been recently executed, or which might not contain any executable code.

The same might be true for Thumb instructions at half-word boundaries, in which case PC[1]=1 but DBGPCSR[1]=0, or ThumbEE instructions at word boundaries, with PC[1]=0 and DBGPCSR[1]=1.

In Cortex-A9, because the DBGPCSR is always a branch target (= start of a basic block to the tool), the debugger should be able to spot many of these cases and attribute the sample to the right basic block.

Workaround

The debugger tools can find the expected PC value and instruction state by reading the DBGPCSR register, and consider it as described in the Description section.

2.5.14 (751480) Conditional failed LDREXcc can set the exclusive monitor

Status

Affects: Product Cortex-A9, Cortex-A9 MPCore.

Fault Type: Programmer Category C

Fault Status: Present in: All r0, r1 and r2 revisions. Fixed in r3p0

Description

A conditional LDREX might set the internal exclusive monitor of the Cortex-A9 even when its condition fails.

Implications

The implication of the erratum is that a subsequent STREX might succeed when it should not. Therefore, the memory region protected by the exclusive mechanism can be corrupted if another agent accesses it at the same time.

Workaround

The workaround for this erratum is to not use conditional LDREX together with non-conditional STREX.

- The erratum cannot trigger unless using conditional LDREX.
- If using conditional LDREX, the associated STREX should also be conditional using the same condition. This means that, even if the exclusive monitor is set by the condition failed LDREX, the following STREX will not execute because it will be condition failed too. For most situations this will naturally be the case anyway.

2.5.15 (752519) An imprecise abort might be reported twice on non-cacheable reads

Status

Affects: Product Cortex-A9, Cortex-A9 MPCore.

Fault Type: Programmer Category C

Fault Status: Present in: All r2, r3 and r4 revisions. Open

Description

In the case when two outstanding read memory requests to device or non-cacheable normal memory regions are issued by the Cortex-A9, and the first one receives an imprecise external abort, then the second access might falsely report an imprecise external abort.

Conditions

The erratum can only happen in systems which can generate imprecise external aborts on device or non-cacheable normal memory regions accesses.

Implications

When the erratum occurs, a second, spurious imprecise abort might be reported to the core when it should not.

In practice, the failure is unlikely to cause any significant issues to the system because imprecise aborts are usually unrecoverable failures. Because the spurious abort can only happen following a first imprecise abort, either the first abort is ignored - and the spurious abort is then ignored too -, or it is acknowledged and probably generates a critical failure in the system, such as a processor reset or whole system reboot.

Workaround

There is no practical software workaround for the erratum.

2.5.16 (754323) Repeated Store in the same cache line might delay the visibility of the Store

Status

Affects: Product Cortex-A9, Cortex-A9 MPCore.

Fault Type: Programmer Category C

Fault Status: Present in: All r2, r3 and r4 revisions. Open

Description

Since r2p0 revision, the Cortex-A9 implements a small counter which ensures the external visibility of all stores in a finite amount of time, causing an eventual drain of the Merging Store Buffer. This is to avoid erratum #754327, where written data could potentially remain indefinitely in the Store Buffer.

This Store Buffer has merging capabilities, and continues to merge data as long as the write accesses are performed in the same cache line. The issue which causes this erratum is that the draining counter resets each time a new data merge is performed.

In the case when a code sequence loops, and continues to write data in this same cache line, then the external visibility of the written data might not be ensured.

A livelock situation might consequently occur if any external agent is relying on the visibility of the written data, and where the writing processor cannot be interrupted while doing its writing loop.

Conditions

The erratum can only happen on Normal Memory regions.

The following examples describe scenarios that might trigger the erratum:

1. The processor continues incrementing a counter, writing the same word at the same address. The external agent (possibly another processor) polls on this address, waiting for any update of the counter value to proceed.
The Store Buffer continues merging the updated value of the counter in its cache line, so that the external agent never sees any updated value, possibly leading to livelock.
2. The processor writes a value in a given word to indicate completion of its task, then continues writing data in an adjacent word in the same cache line.
The external agent continues to poll the first word memory location to check when the processor completes its task. The situation is the same in the first example, because the cache line might remain indefinitely in the merging Store Buffer, creating a possible livelock in the system.

Implications

This erratum might create performance issues, or a worst case livelock scenario, if the external agent relies on the automatic visibility of the written data in a finite amount of time.

Workaround

The recommended workaround for this erratum is to insert a DMB operation after the faulty write operation in code sequences that this erratum might affect, to ensure the visibility of the written data to any external agent.

2.5.17 (756421) Sticky Pipeline Advance bit cannot be cleared from debug APB accesses

Status

Affects: Product Cortex-A9, Cortex-A9 MPCore.

Fault Type: Programmer Category C

Fault Status: Present in: All revisions. Open

Description

The Sticky Pipeline Advance bit is bit[25] of the DBGDSCR register. This bit enables the debugger to detect whether the processor is idle. This bit is set to 1 every time the processor pipeline retires one instruction.

A write to DBGDRCR[3] clears this bit.

The erratum is that the Cortex-A9 does not implement any debug APB access to DBGDRCR[3].

Implications

Because of the erratum, the external debugger cannot clear the Sticky Pipeline Advance bit in the DBGDSCR. In practice, this makes the Sticky Pipeline Advance bit concept unusable on Cortex-A9 processors.

Workaround

There is no practical workaround for this erratum.

The only possible way to reset the Sticky Pipeline Advance bit is to assert the nDBGRESET input pin on the processor, which obviously has the side effect to reset all debug resources in the concerned processor, and any other additional Coresight components nDBGRESET connects to.

2.5.18 (757119) Some *Unallocated memory hint* instructions generate an Undefined Instruction exception instead of being treated as NOP

Status

Affects: Product Cortex-A9, Cortex-A9 MPCore.

Fault Type: Programmer Category C

Fault Status: Present in: All revisions. Open

Description

The ARM Architecture specifies that ARM opcodes of the form 11110 100x001 xxxx xxxx xxxx xxxx are *Unallocated memory hint (treat as NOP)* if the core supports the MP extensions, as the Cortex-A9 does.

The errata is that the Cortex-A9 generates an Undefined Instruction exception when bits [15:12] of the instruction encoding are different from 4'b1111, instead of treating the instruction as a NOP.

Implications

Because of the erratum, an unexpected Undefined Instruction exception might be generated.

In practice, this erratum is unlikely to cause any significant issue because such instruction encodings are not supposed to be generated by any compiler, nor used by any handcrafted program.

Workaround

The workaround for this erratum is to modify the instruction encoding with bits[15:12]=4'b1111, so that the Cortex-A9 treats the instruction properly as a NOP.

If it is not possible to modify the instruction encoding as described, the Undefined Instruction exception handler has to cope with this case, and emulate the expected behavior of the instruction, that is, it must do nothing (NOP), before returning to normal program execution.

2.5.19 (761321) MRC and MCR are not counted in event 0x68

Status

Affects: Product Cortex-A9, Cortex-A9 MPCore.

Fault Type: Programmer Category C

Fault Status: Present in: All revisions. Open

Description

Event 0x68 counts the total number of instructions passing through the register rename pipeline stage. The erratum is that MRC and MCR instructions are not counted in this event.

The event is also reported externally on PMUEVENT[9:8], which suffers from the same defect.

Implications

The implication of this erratum is that the values of event 0x68 and PMUEVENT[9:8] are imprecise, omitting the number of MCR and MRC instructions. The inaccuracy of the total count depends on the rate of MRC and MCR instructions in the code.

Workaround

No workaround is possible to achieve the required functionality of counting precisely how many instructions are passing through the register rename pipeline stage when the code contains some MRC or MCR instructions.

2.5.20 (764319) Read accesses to DBGPRSR and DBGOSLSR may generate an unexpected Undefined Instruction exception

Status

Affects: Product Cortex-A9, Cortex-A9 MPCore.

Fault Type: Programmer Category C

Fault Status: Present in: All revisions. Open

Description

CP14 read accesses to the DBGPRSR and DBGOSLSR registers generate an unexpected Undefined Instruction exception when the DBGSWENABLE external pin is set to 0, even when the CP14 accesses are performed from a privileged mode.

Implications

Because of the erratum, the DBGPRSR and DBGOSLSR registers are not accessible when DBGSWENABLE=0.

This is unlikely to cause any significant issue in Cortex-A9 based systems because these accesses are mainly intended to be used as part of debug over powerdown sequences, and the Cortex-A9 does not support this feature.

Workaround

The workaround for this erratum is to temporarily set the DBGSWENABLE bit to 1 so that the DBGPRSR and DBGOSLSR registers can be accessed as expected.

There is no other workaround for this erratum.

2.5.21 (771221) PLD instructions might allocate data in the Data Cache regardless of the Cache Enable bit value

Status

Affects: Product Cortex-A9, Cortex-A9 MPCore.

Fault Type: Category C

Fault Status: Present in: All r0, r1, r2 and r3 revisions. Fixed in r4p0

Description

PLD instructions prefetch and allocate any data marked as Write-Back (either Write-Allocate or Non-Write-Allocate, Shared or Non-Shared), regardless of the processor configuration settings, including the Data Cache Enable bit value.

Implications

Because of this erratum, unexpected memory cacheability aliasing is created which might result in various data consistency issues.

In practice, this erratum is unlikely to cause any significant issue. The Data Cache is likely to be enabled as soon as possible in most systems, and not dynamically modified. Therefore, this erratum is likely to impact only boot-up code. This code is usually carefully controlled and does not usually contain any PLD instruction while Data Cache is not enabled.

Workaround

If this erratum impacts a system, a software workaround is available which is to set bit [20] in the undocumented Control register, which is placed in CP15 c15 0 c0 1.

This bit needs to be written with the following Read/Modify/Write code sequence:

```
MRC p15,0,r0,c15,c0,1
```

```
ORR r0,r0,#0x00100000
```

```
MCR p15,0,r0,c15,c0,1
```

Setting this bit causes all PLD instructions to be treated as NOPs, with the consequence that code sequences that usually use the PLDs, such as the memcpy() routine, might suffer from a visible performance drop. Therefore, if this workaround is applied, ARM strongly recommends restricting its use to periods of time where the Data Cache is disabled.

2.5.22 (771224) Visibility of Debug Enable access rights to enable/disable tracing is not ensured by an ISB

Status

Affects: Product Cortex-A9, Cortex-A9 MPCore.

Fault Type: Programmer Category C

Fault Status: Present in: All revisions. Open

Description

According to the ARM architecture, any change in the Authentication Status Register should be made visible to the processor after an exception entry or return, or an ISB.

Although this is correctly achieved for all debug-related features, the ISB is not sufficient to make the changes visible to the trace flow. As a consequence, the **WPTTRACEPROHIBITEDn** signal(s) remain stuck to their old value up to the next exception entry or return, or to the next serial branch, even when an ISB executes.

A serial branch is one of the following:

- Data processing to PC with the S bit set (for example, MOVS pc, r14)
- LDM pc ^

Implications

Because of the erratum, the trace flow might not start, nor stop, as expected by the program.

Workaround

To work around the erratum, the ISB must be replaced by one of the events causing the change to be visible. In particular, replacing the ISB by a MOVS PC to the next instruction will achieve the correct functionality.

2.5.23 (771225) Speculative cacheable reads to aborting memory regions clear the internal exclusive monitor, may lead to livelock

Status

Affects: Product Cortex-A9, Cortex-A9 MPCore.

Fault Type: Programmer Category C

Fault Status: Present in: All r2 and r3 revisions. Fixed in r4p0

Description

On Cortex-A9, when a cacheable read receives an external abort, the aborted line is allocated as invalid in the Data Cache, and any allocation in the Data Cache clears the internal exclusive monitor.

Therefore, if a program executes a LDREX/STREX loop which continues to receive an abort answer in the middle of the LDREX/STREX sequence, then the LDREX/STREX sequence never succeeds, leading to a possible processor livelock.

As an example, the following code sequence might exhibit the erratum:

```
loop LDREX
...
DSB
STREX
CMP
BNE loop
....
LDR (into aborting region)
```

The LDREX/STREX does not succeed on the first pass of the loop, and the BNE is mispredicted, and the LDR afterwards is speculatively executed.

Therefore the processor keeps on executing:

```
LDR to aborting region (this speculative LDR now appears "before" the LDREX & DSB)
LDREX
DSB
STREX
```

The LDR misses in L1, and never gets allocated as valid because it is aborting

The LDREX executes, and sets the exclusive monitor

The DSB executes. It waits for the LDR to complete, which aborts, causing an allocation (as invalid) in the Data Cache, which clears the exclusive monitor.

The STREX executes, but the exclusive monitor is now cleared, so the STREX fails.

The BNE might be mispredicted again, therefore the LDR is speculatively executed again, and the code loops back on the same failing LDREX/STREX sequence.

Conditions

The erratum happens in systems which might generate external aborts in answer to cacheable memory requests.

Implications

If the program reaches a stable state where the internal exclusive monitor continues to be cleared in the middle of the LDREX/STREX sequence, then the processor might encounter a livelock situation.

In practice, this scenario is very unlikely to happen because several conditions might prevent it:

- Normal LDREX/STREX code sequences do not contain any DSB, so it is very unlikely that the system would return the abort answer precisely in the middle of the LDREX/STREX sequence on each iteration.
- Some external irritators (for example, interrupts) might happen and cause timing changes which might exit the processor from its livelock situation.
- Branch prediction is usually enabled, so the final branch in the loop is usually predicted correctly after a few iterations of the loop, preventing the speculative LDR from being issued, so that the next iteration of the LDREX/STREX sequence succeeds.

Workaround

For this erratum, either of the following workarounds fix the problem:

- Turn on the branch prediction.
- Remove the DSB in the middle of the LDREX/STREX sequence. If a DSB is required, ARM recommends that you place it before the LDREX/STREX sequence, and implement the LDREX/STREX sequence as recommended by the ARM architecture.

2.5.24 (775419) PMU event 0x0A (exception return) might count twice the LDM PC ^ instructions with base address register write-back

Status

Affects: product Cortex-A9, Cortex-A9 MPCore.

Fault Type: Programmer Category C

Fault Status: Present in: All revisions. Open

Description

The LDM PC ^ instructions with base address register write-back might be counted twice in the PMU event 0x0A, which is counting the number of exception returns.

The associated PMUEVENT[11] signal is also affected by this erratum, and might be asserted twice by a single LDM PC ^ with base address register write-back.

Implications

Because of the erratum, the count of exception returns is imprecise. The error rate depends on the ratio between exception returns of the form LDM PC ^ with base address register write-back and the total number of exceptions returns.

Workaround

There is no workaround to this erratum.

2.5.25 (782774) A spurious event 0x63, STREX passed, can be reported on an LDREX that is preceded by a write to Strongly Ordered memory region

Status

Affects: product Cortex-A9, Cortex-A9 MPCore.

Fault Type: Programmer Category C

Fault Status: Present in: All r2 and r3 revisions. Fixed in r4p0

Description

A write to Strongly Ordered memory region, followed by the execution of an LDREX instruction, can cause the *STREX passed* event to be signaled even if no STREX instruction is executed.

As a result, the event 0x63 count might be faulty, reporting too many *STREX passed* events.

This erratum also affects the associated PMUEVENT[27] signal. This signal will report the same spurious events.

Conditions

The erratum requires the following conditions:

1. The processor executes a write instruction to a Strongly Ordered memory region.
2. The processor executes an LDREX instruction.
3. No DSB instruction is executed, and there is no exception call or exception return, between the write and the STREX instructions.

Under these conditions, if the write instruction to Strongly Ordered memory region receives its acknowledge (BRESP response on AXI) while the LDREX is being executed, the erratum can happen.

Implications

The erratum leads to a faulty count of event 0x63, or incorrect signaling of PMUEVENT[27].

Workaround

The workaround for this erratum is to insert a DMB or DSB instruction between the write to Strongly Ordered memory region and the LDREX instruction.

